Biases in allocation under risk and uncertainty: Partition dependence, unit dependence, and procedure dependence

Thomas Langer
University of Muenster

Craig R. Fox
University of California at Los Angeles

Abstract

Previous studies of employee investment in retirement plans suggest that people typically "naively diversify" their investment funds, tending to allocate 1/*n* of the total to each of *n* available instruments (Benartzi & Thaler, 2001). In this paper we provide experimental evidence that this bias extends to allocation among simple chance prospects and demonstrate three new violations of rational choice theory implied by use of the naïve diversification strategy. Study 1 demonstrates "partition dependence" in which participants' allocations among a fixed set of investments varies with the hierarchical structure of the option set (e.g., by vendor and instrument). Study 2 demonstrates "unit dependence" in which participants' preferred allocations vary with the metric in which the investment is reported (dollars versus number of shares). Study 3 demonstrates "procedure dependence" in which the bias toward even allocation disappears if participants are asked to choose from a menu of possible portfolios. We show that these results extend to sophisticated participants, simple well-specified gambles and incentive-compatible payoffs. We close with a discussion of theoretical and prescriptive implications.

Running Head: Dependencies in Naïve Diversification

Address Correspondence to:
Thomas Langer
University of Muenster
Universitaetsstr. 14-16
48143 Muenster, Germany
Thomas.Langer@wiwi.uni-muenster.de

The authors gratefully acknowledge financial support from the Deutsche Forschungsgemeinschaft (DFG) and the SFB 504 at the University of Mannheim. The present work was also supported by grant SES-00-99209 from the National Science Foundation to the second author. Finally, we acknowledge the support of the Fuqua School of Business, Duke University, who supported the authors when these studies were initially conceived. This paper has benefited from discussions with Shlomo Benartzi. We are grateful for helpful comments on previous drafts of this paper from David Bardolet and Brett Rogers.

Biases in allocation under risk and uncertainty:
Partition dependence, unit dependence, and procedure dependence

Organizations and individuals must often decide how to allocate resources among projects whose consequences are risky or uncertain. Firms must periodically choose how to divide capital and human resources among various research and development projects; venture capitalists must decide how to distribute funds among start-up companies; law firms must decide how they will invest time and money among a set of cases whose payout is uncertain. Among the most important financial decisions that individuals make is how to allocate their savings among potential investments. In the last few decades, American employers have rapidly shifted from offering defined retirement benefit plans to offering defined contribution retirement saving plans in which employees choose how to allocate their money among a fixed set of possibilities. As of 1999 about 85 percent of private retirement savings contributions went to accounts in which individuals decided how to invest plan assets (Poterba, Venti & Wise, 2001), and as of the end of 2003 there were an estimated 42 million participants and $1.9 trillion in total assets invested in 401k plans (Employer Benefit Research Institute, 2004).

In light of the importance and prevalence of risky and uncertain allocation, there should be great interest in understanding how people make such decisions and how they might be improved. To date, however, there has been relatively little behavioral research on how people allocate funds over a fixed set of risky or uncertain prospects. Especially noteworthy is a groundbreaking investigation of personal savings decisions in which Benartzi and Thaler (2001; hereafter “BT”) argued that people typically employ “naïve diversification” strategies, allocating $1/n$ of their funds to each of n investment prospects available to them with little regard to the nature of these investments. For instance, in a survey of UCLA employees, these investigators found that a substantial proportion of participants allocated half of their (hypothetical) retirement savings to each of two investments with little regard to whether the two options were a stock fund and a bond fund, a stock fund and a balanced stock/bond fund, or a bond fund and a balanced stock/bond fund. In another survey UCLA employees offered four fixed-income funds and one equity fund allocated a median 40% of their retirement savings to equities, whereas employees offered one fixed-income fund and four equity funds allocated a median 75% of their retirement savings to equities. BT found further evidence of naïve diversification in their analysis of actual investments in 170 retirement savings plans obtained in a database from Money Market Directories. The results showed that employees allocated more of their savings to equities if their company offered a greater proportion of equity investment options. For instance, University of California employees, who allocated their retirement savings among a stock fund and four bond funds, invested 34 percent of their money in stocks, whereas TWA employees, who allocated their retirement savings among five stock funds and one bond fund, invested 75 percent of their money in stocks.

BT's results suggest that naïve diversification is a strong and robust phenomenon in individual decisions concerning retirement investment, but it remains to be seen to what extent this phenomenon would be observed in other domains of decision under uncertainty.[1] For instance, it could be that naïve diversification is a prescriptive norm that people apply only to personal investment decisions. Or it could be that people only apply naïve diversification in situations where the probability distribution over outcomes or the correlation among returns of potential investments is unclear. The first purpose of the present paper is to investigate whether naïve diversification is exhibited by sophisticated respondents making allocations among a small number of simple well-defined lotteries with incentive-compatible payoffs and transparent correlations.

The second and more central purpose of this paper is to characterize and test new implications of naïve diversification. In particular we conjecture that this strategy will give rise to three systematic dependencies that cannot easily be reconciled with rational choice theory and pose a substantial challenge for decision analysis. First, if decision makers rely on naive diversification to even a modest extent, the final distribution will be affected by the particular grouping of options over which resources are allocated. For instance, participants in retirement plans who are offered a choice of vendors followed by a choice of

---

[1] BT did note that naive diversification may be an instance of a more general *diversification heuristic* that also includes the tendency to seek variety when making multiple choices among snacks and other consumer products (e.g., Simonson, 1990; Read & Loewenstein, 1995; Fox, Ratner & Lieb, 2005), but we are aware of no studies of naïve diversification in allocation among risky or uncertain prospects outside of an investment context.

instruments within vendors (as is the case with the typical 403b plan offered by not-for-profit organizations) might first allocate evenly across vendors, then evenly across instruments within vendors. Participants in a retirement plan who are offered these same instruments in parallel (as is the case with the typical 401k plan offered by for-profit organizations) would instead allocate evenly across instruments with little regard to their corresponding vendor.[2] We refer to the sensitivity of allocations to the way in which prospects happen to be grouped as *partition dependence*.

Second, if decision makers are biased toward allocating resources evenly over potential investments, the final distribution will depend on what, specifically, is being allocated (i.e., the unit of measurement). For instance, brokerage accounts typically require investors to indicate the number of *shares* of common stocks that they wish to purchase, whereas they typically require investors to indicate the number of *dollars* worth of mutual funds that they wish to purchase. Thus, if share prices of investment instruments differ then the strategy of allocating money evenly among instruments will yield a different portfolio than the strategy of purchasing an equal number of shares of each instrument. To illustrate, suppose a person wishes to invest $3,000 in two stocks: stock A, currently trading at $10 per share, and stock B, currently trading at $20 per share. In this case, a purchase of 100 shares of stock A and 100 shares of stock B may seem more attractive than a purchase of 150 shares of stock A and 75 shares of stock B. However, the opposite may be true when the same portfolios are expressed in terms of dollars invested: $1000 worth of stock A and $2000 worth of stock B may seem less attractive than $1500 worth of stock A and $1500 worth of stock B. We refer to the tendency for portfolio choice to be affected by the particular unit in which allocations are expressed as *unit dependence*.

Finally, if people rely on naïve diversification when allocating funds among a fixed set of prospects, the question arises to what extent they will do so when choosing among possible mixtures. Previous work in judgment and decision making has found that response strategies are influenced by the compatibility between the strategy and the response mode (Slovic, Griffin & Tversky, 1990; Fischer &

[2] Indeed, when the second author originally signed up for the 403(b) plan at his first tenure-track job, he asked the university benefits counselor how most people allocated among the three available vendors, to which the counselor responded, “most people allocate 1/3 to each.”

Hawkins, 1993). In this case we predict that the proportion of resources invested in each prospect should be more salient to decision makers when they are called on to make an explicit allocation than when they are asked to choose between prespecified portfolios. Moreover, the heuristic to distribute equally will be more accessible when allocating whereas other choice heuristics such as maximizing the minimum possible return or maximizing expected value will be more accessible when choosing a portfolio, especially if this information is made more transparent. Thus we predict that explicit allocation will lead to more even distributions than an equivalent choice among prespecified portfolios that are mixtures of the same underlying assets. We refer to the tendency of different elicitation procedures to give rise to different portfolio choices as *procedure dependence.*

In this paper we provide an experimental investigation of naïve diversification in allocation under risk and uncertainty that extends previous work by BT. First, we seek to replicate the basic result that people are biased toward allocating 1/*n* of their budget to each of *n* investment prospects, with insufficient regard to the nature of those prospects. We depart from previous work by using not only simple investments but also well-defined chance lotteries with incentive-compatible payoffs. Second, we explore the extent to which participants will exhibit the three hypothesized dependencies under such conditions. In Study 1 we seek evidence of partition dependence for both hypothetical investments and chance lotteries. In Study 2 we test for unit dependence for both hypothetical investments and chance lotteries. In Study 3 we seek evidence of procedure dependence in a within-subject experiment involving chance lotteries.

Study 1: Partition Dependence

In this study we investigate whether the allocation of funds is affected by the grouping of lotteries or investment instruments. As mentioned earlier, many retirement savings plans call on employees to first allocate funds by vendor then by specific instruments within each vendor. Many firms offer a focal investment such as company stock versus a set of other instruments, in which case employees may allocate first between this focal investment and other investments, then proceed to allocate among remaining investments.

Preliminary evidence of partition dependence can be found in an analysis of field data reported in BT. They find that among a wide range of pension plans that do not include the option of company stock, employees allocate their retirement savings roughly evenly among stocks and bonds (49 percent to equities), whereas among pension plans that do offer company stock as an option, employees invest 42 percent in company stock, 29 percent in other equities, and 29 percent in fixed-income investments. This pattern could arise from hierarchical partitioning in which participants who are offered company stock first divide their savings roughly equally between company stock and all other instruments, then divide equally among remaining instruments.

We seek more direct evidence of partition dependence in two studies in which we hold a set of risky or uncertain prospects constant, but vary the subjective grouping of these prospects in various ways. In Study 1A we test for partition dependence using three simple generic investment instruments that are assigned to one of two vendors; participants are first asked to allocate funds among vendors then among instruments offered by a common vendor. In Study 1B we explore whether partition dependence might be observed in the case of simple chance lotteries with incentive-compatible payoffs.

Study 1A

Method

We recruited 184 MBA students at Duke University during an orientation session to complete a brief survey in exchange for a donation to charity. Participants were asked to suppose that they had an opportunity to invest a portion of their income, with a matching contribution from their employer, in tax-deferred investments for retirement. Next, they were asked how they would divide this benefit among three asset classes: bonds (long-term U.S. treasury notes), equities (an S&P 500 stock index fund), and real estate (a nationally diversified Real Estate Investment Trust). Assets were grouped so that two were offered by one fictitious vendor ("Fiduciary") and one was offered by a second fictitious vendor ("Integrity"). The mapping from vendor to asset was fully counterbalanced (so that there were three meaningful variations, one with each asset mapped to the singleton vendor) and the order of presentation was also counterbalanced. After all of the vendors and assets were described, the allocation decision was presented as a two-stage process. In the first stage, participants were asked to allocate their endowment (in percentage terms) among the two vendors. In the second stage, participants were asked to subdivide the money allocated to the two-asset-vendor among its two assets. For one group of participants ($n = 61$) the one-asset vendor offered bonds (treatment B); for a second group of participants ($n = 63$) the one-asset vendor offered stocks (treatment S); for a third group of participants ($n = 60$) the one-asset vendor offered real estate (treatment R).

We hypothesized that people would be biased toward even allocation among the available vendors, and also biased toward even allocation among assets within vendors. Thus, participants who employ perfect naïve diversification should allocate half of their savings to each vendor, and half to each instrument within the vendor that offers two instruments. For example, in treatment R perfect naïve diversifiers would allocate 50% to real estate, and 25% each to stocks and bonds. More generally, we expected a bias in this direction so that allocations are biased toward 50% for singleton assets and 25% for non-singleton assets.

Hypothesis 1A: Each asset will be allocated more funds when it is the only asset assigned to a vendor than when it is one of two assets assigned to a vendor.

Results

Table 1 displays mean and median allocations to the three assets for the treatments B, S and R; means are also depicted visually in Figure 1. For instance, the mean allocation into bonds was 35% when bonds were the singleton (treatment B), compared to a 21% allocation, on average, when paired with another instrument (treatments S and R, see Table 1). The difference between the bond allocation in treatment B and the two other treatments is highly significant ($p<0.0001$ by Mann-Whitney). Even higher differences are observed in the respective analysis of stocks (59% in treatment S vs. 38% in the other treatments) and real estate (40% in treatment R vs. 23% in the other treatments). These differences are also highly significant ($p<0.0001$).[3]

[3] We further asked subjects to rate their own knowledge concerning such investments on a scale from 0 (no knowledge) to 10 (a great deal of knowledge) and found that self-rated knowledge does not influence the strength of the effect. The only significant effect of knowledge is a generally higher allocation into stocks.

It is worth pointing out that although participants were strongly influenced by the grouping of investments, they apparently did distinguish among them: Median allocations to stocks were above the proportions implied by pure two-step naïve diversification in all three experimental conditions (i.e., above 25%, 50%, and 25% in treatments, B, S, and R, respectively) and median allocations to bonds and real estate were below proportions implied by pure naïve diversification in all three conditions.

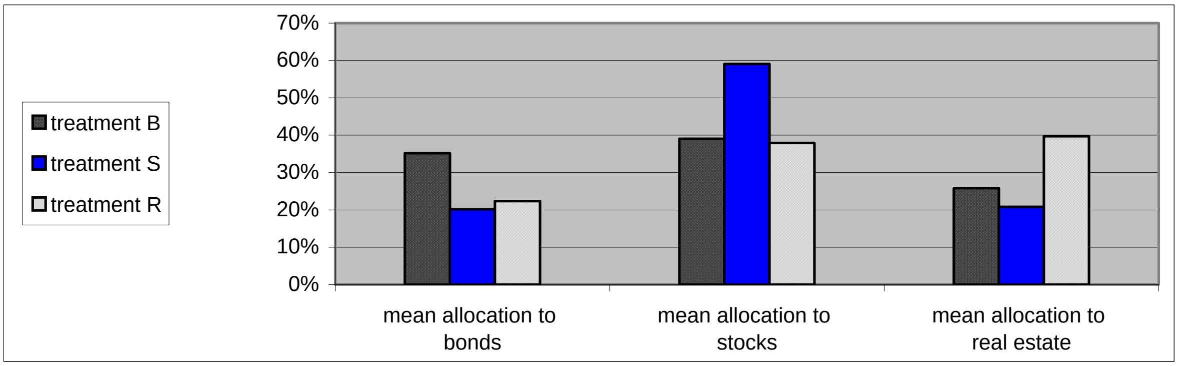

Figure 1: Results of Study 1A.

		% allocation to bonds	% allocation to stocks	% allocation to real estate
		mean *(median)*	mean *(median)*	mean *(median)*
treatment B	n=61	35.2 *(30.0)*	39.0 *(35.0)*	25.8 *(24.0)*
treatment S	n=63	20.1 *(15.0)*	59.1 *(65.0)*	20.8 *(20.1)*
treatment R	n=60	22.4 *(20.0)*	37.9 *(38.3)*	39.7 *(37.5)*

Table 1: Mean and median allocations to bonds, stocks and real estate in Study 1A.

Study 1B

We wished to test the robustness of partition dependence through a replication. Study 1B differed from Study 1A in three important respects. First, participants were asked to allocate their endowment among chance lotteries with explicit probability distributions over outcomes rather than more abstract assets with more ambiguous (i.e. vaguer) distributions. Second, we introduced incentive-compatible payments. Third, in order to obtain a neutral benchmark of allocation preferences we included a non-hierarchical treatment in which assets were not grouped by vendors.

Method

We recruited 171 Masters degree students specializing in Banking and Finance at the University of Mannheim. Participants received a questionnaire in which they were asked to divide an endowment of €30,000 between three assets. These assets were presented as simple two-outcome lotteries and had the following form:

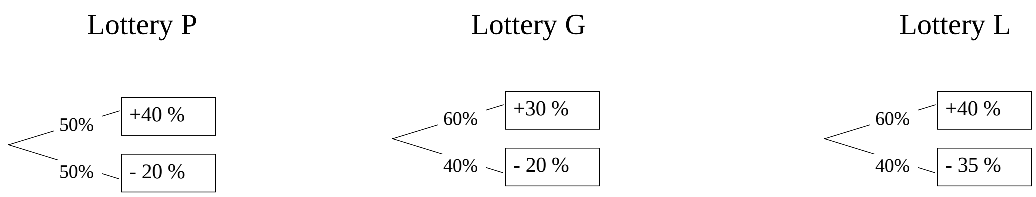

Although all three lotteries have the same expected return (10%), each is distinguishable from the other two lotteries by some specific property that we label here for expository purposes P, G, and L. Lottery *P* has a lower gain *probability* than *G* and *L*; lottery *G* has a lower possible *gain* than *P* and *L*; lottery *L* has a higher possible *loss* than *P* and *G*. (In the actual survey, the three lotteries were labelled “Asset 1,” “Asset 2,” and “Asset 3” in the order in which they were described).

In the *neutral* condition, participants (*n* = 59) were asked to make a single decision how to allocate their endowment among assets *P*, *G* and *L*. In the three hierarchical treatments, lotteries were assigned to vendors A and B. As in Study 1A, one vendor offered a single asset and one vendor offered two assets. The distinguishing property of the lotteries was mentioned to explain the assignment of assets to vendors. For example, in condition *P* (*n* = 38), it was noted that vendor A offers two assets (*G* and *L*) both with a gain probability of 60%, while vendor B offers only one asset (*P*) with the lower gain probability of 50%. In the same way, assets *G* and *L* were singled out in the conditions *G* (*n* = 40) and *L* (*n* = 34), respectively. The order in which assets and vendors were presented was counterbalanced.

Participants were informed that five respondents would be selected at random to be paid real money on the basis of their choices. For “real money” participants, the outcome of the lotteries would be determined in front of them and the final value of the portfolio would be paid to those participants immediately (based on a starting endowment of €30 instead of €30,000). Actual payments ranged from €20.85 to €42.00 with a mean of €35.09.

We predicted that participants would invest more into assets when they are assigned to a one-asset vendor (in which case there should be a bias toward 1/2) than when they are assigned to a two-asset vendor (in which case there should be a bias toward 1/4), and that the allocation to each asset should be between these values in the neutral condition (in which case there should be a bias toward 1/3). We note that because our procedure highlighted a negative property of the singleton in all cases (lower gain probability of *P*, lower gain size of *G*, higher loss size of *L*), this should, if anything, discourage investment in singletons and weaken our hypothesized effect.

Hypothesis 1B: Assets will receive greater investment when assigned to one-asset vendors than when assigned to two-asset vendors. Assets should receive an intermediate investment in the neutral condition.

Results

The results of Study 1B are displayed in Figure 2 and summarized in Table 2. As predicted, each asset received a significantly higher allocation in the treatment where it was the singleton ($p<0.01$, Mann-Whitney one-tailed in all three cases). Interestingly, in conditions *P* and *G* the median allocation to the singleton was precisely 50%. For all assets, the median allocation in the neutral treatment is between the median allocations in the singleton and nonsingleton treatments. The difference between the median allocation to the singleton in the hierarchical treatments and the neutral treatment was 10% (=50% - 40%) for treatment P, 17% for treatment G, and 10% for treatment L. By a one-tailed Mann-Whitney test this difference is highly significant for treatment G ($p<0.01$) and significant for treatments P and L ($p<0.05$).

It is worth observing as in Study 1A that although participants exhibited strong (partition dependent) naïve diversification, they also apparently distinguished among the lotteries, favoring Asset P to Asset G to Asset L. The design of Study 1B allowed us to examine the relative contribution of (partition dependent) naïve diversification and “neutral” preferences derived from treatment N. Define the “ignorance prior” proportion (IP) as the proportion of funds implied by even allocation (i.e., 50% when an asset is in the role of the singleton and 25% when an asset is paired with another asset) and “neutral” proportion (NP) as the mean proportion allocated to that asset in the neutral condition. We regressed mean proportion allocated (PA) to each asset on the ignorance prior proportion and the neutral proportion (with no intercept term):

$$PA = \beta_1\ IP + \beta_2\ NP$$

We obtain an excellent fit of the model, $F(2, 7) = 220.84$, $p = .0000$; adjusted $R^2 = .98$.[4] Moreover, the ignorance prior and neutral proportions both received significant and roughly equal weights that summed to near unity ($\beta_1 = 0.53$, SE = 0.12, $p = .002$; $\beta_2 = 0.48$, SE = 0.12, $p = .005$). Thus, allocations appear to reflect both a tendency to naively diversify in a partition dependent manner and also a tendency to distinguish among assets. Of course, it is worth noting that allocations in the "neutral" treatment are presumably biased toward the ignorance prior, so we surmise that the beta weights in the foregoing regression actually underestimate relative reliance on naïve diversification exhibited by our participants.

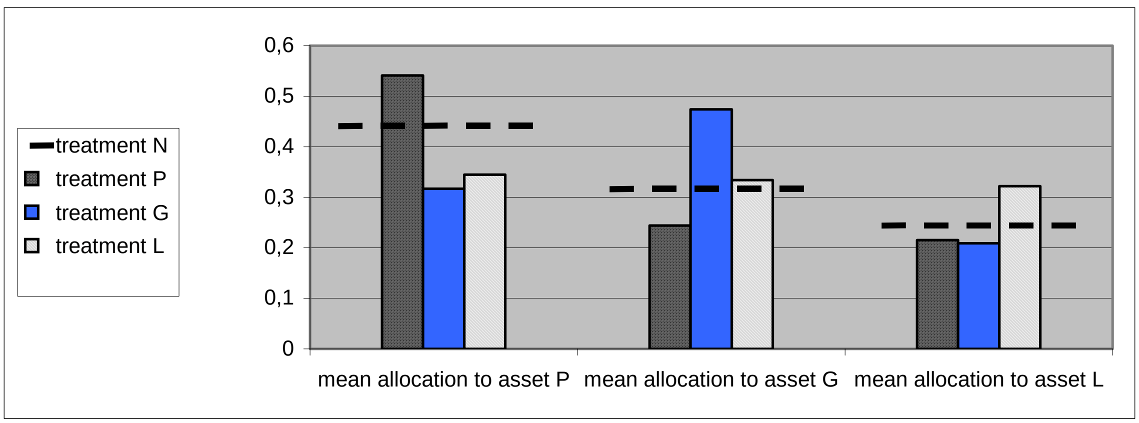

Figure 2: Results of Study 1B.

		% allocation to Asset P	% allocation to Asset G	% allocation to Asset L
		mean (median)	mean (median)	mean (median)
treatment P	*n* = 38	54.1 (50.0)	24.4 (28.0)	21.5 (18.0)
treatment G	*n* = 40	31.7 (23.0)	47.4 (50.0)	20.9 (14.3)
treatment L	*n* = 34	34.5 (32.7)	33.4 (25.5)	32.2 (30.0)
treatment N	*n* = 59	**44.3 (40.0)**	**31.6 (33.0)**	**24.1 (20.0)**

Table 2: Mean and median allocations to assets P,G and L in Study 1B.

[4] Of course the allocation to asset L is determined by the allocation to assets P and G. This dependency does not result in a systematic overestimation of the fit of the model though. If e.g. we completely ignored the allocations to asset L in our analysis, the fit of the model would even be better (adj. $R^2 > 0.99$).

Study 2: Unit Dependence

We next turn to the question of unit dependence. In this section, we present the results of two studies in which we hold the set of risky prospects constant, but vary the units that are being allocated (money versus number of shares). In Study 2A we explore this phenomenon using stock investments and hypothetical allocations. In Study 2B we examine whether this phenomenon extends to simple well-defined chance lotteries and incentive-compatible consequences.

Study 2A

Method

We recruited 214 graduate students at Duke University to complete a survey that also included several unrelated items. We asked participants to imagine that they had been given $20,000 by a relative that they could invest in shares of IBM and Apple stocks. All participants were also told the latest closing prices of these stocks and to assume that the stocks could be purchased at these prices ($91.72 for IBM and $15.21 for Apple). They were then asked to choose between a portfolio that entailed an equal number of *shares* invested in each stock (henceforth we refer to this as portfolio *S*) and a portfolio that entailed an equal amount of *money* invested in each stock (henceforth portfolio *M*), both subject to the constraint that only an integer number of shares could be purchased. As illustrated in Figure 3, these portfolios were presented either in terms of the amount of money invested in each stock (condition *M*) or in terms of the number of shares purchased of each stock (condition *S*). On the actual survey we used the labels "Portfolio 1" and "Portfolio 2" and participants were not explicitly told the rationale for constructing these portfolios. The order of presentation was counterbalanced.

Condition M:

Portfolio M:	Portfolio S:
$ 9,992.97 worth of Apple	$ 2,844.27 worth of Apple
$ 9,997.48 worth of IBM	$ 17,151.64 worth of IBM
$ 9.55 cash	$ 4.09 cash
$ 20,000.00 total	$ 20,000.00 total

Condition S:

Portfolio M:	Portfolio S:
657 shares of Apple	187 shares of Apple
109 shares of IBM	187 shares of IBM
$ 9.55 cash	$ 4.09 cash
= $ 20,000.00 total	= $ 20,000.00 total

Figure 3: Different presentation modes for portfolio M and portfolio S.

We predicted that the unit used to describe portfolios would be more accessible than the alternative unit so that respondents would be more likely to favor a portfolio that is equal in terms of that unit. Thus, participants should find portfolio *M* more attractive when both portfolios are described in terms of money invested (condition *M*) than when both portfolios are described in terms of shares purchased (condition *S*); the converse should be true for portfolio *S*.

Hypothesis 2A: A larger proportion of respondents will favor portfolio *M* to portfolio *S* in condition *M* than in condition *S*.

Results

Results of Study 2A, displayed in the top half of Table 3, accord with our prediction. In condition *M*, most participants (58%) chose the portfolio *M*, which was an even split in money terms. In condition *S*, most participants (59%) opted for portfolio *S*, the even split in share terms ($\chi^2(1) = 5.99$, $p<0.025$).

Study 2A	Treatment M (money) *n* = 98	Treatment S (shares) *n* = 116
% of participants preferring portfolio M	**58.2**	**41.4**
% of participants preferring portfolio S	**41.8**	**58.6**
	p<0.025	
Study 2B	Treatment M (money) *n*= 67	Treatment S (shares) *n*= 67
Mean (median) % of endowment invested into less expensive asset A	42.4 (40.0)	16.6 (12.5)
	p<0.001	

Table 3: Results of Studies 2A and 2B.

Study 2B

Method

We wished to test the robustness of unit dependence in a replication. Study 2B differs from the previous study in three respects. First, the available assets were described as explicit probability distributions over returns; second, we provided monetary incentives; and third, participants were asked to make their own allocations on a continuous scale rather than choose between two fixed allocations. We recruited 134 undergraduate business students from the University of Mannheim. They were asked to divide an endowment of €56,000 between two assets. Participants were told that the share price of asset A was €100 and the share price of asset B was €700. Both assets could increase or decrease in value with equal probability. In the questionnaires the share prices and potential final values were displayed graphically as follows:

Asset A: 100 € → 50%: $\boxed{170\ €}$; 50%: $\boxed{50\ €}$

Asset B: 700 € → 50%: $\boxed{940\ €}$; 50%: $\boxed{600\ €}$

In the survey we named the assets PRIMAS and SECUNDAS. Name assignment and presentation orders were counterbalanced. Participants were informed that three respondents would be selected at random and paid according to the final value of their portfolio (based on an endowment of €56 instead of €56,000). Lotteries would be resolved using a chance device in front of these "real money" participants. Actual payments ranged from €47 to €73, with a mean of €63.50.

One group of participants (n = 67) was asked to indicate the percentage of the available *money* they would like to invest in each asset (condition M). The second group of participants (n = 67) was asked to indicate the percentage of *shares* they would like to purchase of each asset using the available money (condition S). Just before soliciting responses we reminded participants of the units in which they were designating their responses: "Please note that all of your statements refer to monetary terms [share numbers]" (translated from the original German).

Note that an even split in money terms (€28000, €28000) corresponds to an uneven number of shares (280, 40), whereas an even number of shares (70, 70) corresponds to an uneven allocation of money (€7000, €49000). We thus hypothesized that the tendency toward an even allocation would result in a higher allocation to the less expensive asset A when the allocation was specified in money terms (condition *M*).

Hypothesis 2B: A higher proportion of the endowment will be invested into asset A (the less expensive asset) in condition *M* than in condition *S*.

Results

The results of Study 2B are presented in the lower half of Table 3. All numbers are specified in terms of the percentage of money that participants allocated to asset A. Thus, whereas these numbers correspond to participants' raw responses for condition *M*, they are translated from share percentages indicated by participants in condition *S*. For instance, the reported median value 12.5% in condition *S* reflects an even allocation in share terms (70 shares of each type) as it corresponds to an investment of €7000 (12.5% of €56000) into asset A. These results replicate those of Study 2A and confirm our prediction: participants allocated more money to the less expensive asset A if they specified percentages in terms of money rather than shares (mean percentages 42.40 in condition *M* vs. 16.55 in condition *S*). This difference is highly significant by a Mann-Whitney test ($p<0.001$).

Study 3: Procedure Dependence

We now turn to the question of procedure dependence and hypothesize that portfolios will reflect more even allocation among instruments when they are selected by explicit allocation than when they are selected from a set of possible aggregate portfolios. BT provided indirect evidence for this hypothesis in a hypothetical study in which some participants decided how to allocate savings between a stock fund and a bond fund. A second group of participants chose between five funds that were (unbeknownst to them) mixtures of these two funds. The discrepancy between groups was striking: participants in the "allocation" condition distributed 56 percent of total savings to the stock fund, whereas participants in the "choice" condition distributed 75 percent of total savings to stocks.

Our goal in Study 3 was to replicate this effect in a tightly controlled environment with well-specified gambles and incentive-compatible payoffs. Furthermore, we hoped to demonstrate naïve diversification and procedure dependence among individuals in a within-subject experimental design.

Finally, we intended to isolate the key procedural attributes that give rise to differences in allocations among risky prospects (and thereby eliminate a confound in Benartzi & Thaler's study).

It is important to note that BT's second "choice" condition differed from the "allocation" condition in three ways. First, and most obviously, the "choice" condition required a choice of the most preferred portfolio whereas the "allocation" condition required an explicit allocation among base investments. Second, the "choice" condition presented information concerning the historical distribution of returns of the aggregate portfolio whereas the "allocation" condition only provided distributions of the base investments. Third, the "choice" condition explicitly mentioned the average expected return of each mixture portfolio whereas the "allocation" condition did not. Thus, it may be that BT obtained their result because they provided participants in the "choice" condition with aggregate expected payoffs (thereby making especially salient the dimension on which pure equity portfolios are superior to mixed portfolios). Even if this is not the case, it is an interesting question to what extent the shift toward equities in the "choice" condition is driven by the elicitation *task* (choice) and to what extent it is driven by *information* provided in elicitation (probability distribution of outcomes of the aggregate portfolio). Thus, in Study 3 we ask participants to: (1) allocate funds among two simple chance lotteries, (2) choose among the set of possible portfolio mixtures of these two lotteries, in which we provide participants with the probability distribution over returns (but unlike BT do not explicitly mention average expected returns), and (3) allocate funds among two chance lotteries for which they are also provided with information concerning the probability distribution over returns implied by each allocation that they tentatively select. In this way we can establish the robustness of the procedure dependence phenomenon and also diagnose the extent to which this phenomenon is driven by the task (allocation versus choice) versus the information (whether or not people are provided the probability distribution over outcomes of the aggregate portfolio).

Method

We recruited 52 MBA students from Duke University for a 30-minute decision making task to be completed on a computer. We contributed $5 for each participant to a charity. In addition, participants were told that we would stake one respondent with $250 for sure, plus the return on $1,000 invested according to one of their choices. We told participants that we would announce the "real money" participant, his or her decision, and the result publicly, and donate this sum to charity as well.

Study 3 consisted of three phases in which participants indicated their most preferred portfolio from a continuous set of possibilities. All available portfolios were convex combinations of two risky prospects A and B, though this was not obvious to participants in all experimental conditions. The outcomes of prospects were both to be determined by the single throw of a die. Thus, the prospects A and B were not independent of each other. We took great care to ensure that participants understood how the outcomes of paired prospects were related to one another.[5] We used three different primary base pairs of prospects or "scenarios" (see Figure 4). Consider, for example, Scenario 3. If a roll of a fair die would land three or four, money invested in option A would decrease in value by 24%, whereas money invested in option B would increase by 20%. Scenario 1 was designed with prospects that entailed pure gains, whereas Scenarios 2 and 3 contained mixed gambles that entailed both potential losses and gains such that for any mixture it would be impossible to completely avoid the possibility of a loss.

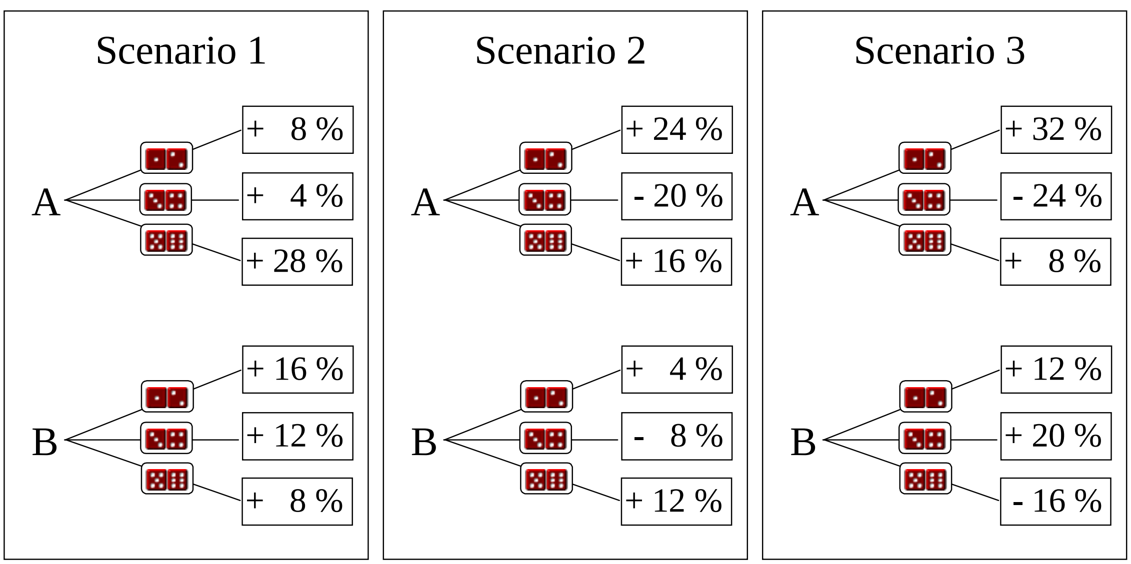

Figure 4: Base prospects used in Study 3.

[5] Participants were required to work through control tasks before they could proceed with the experiment. A typical control question was, "How do you have to divide your endowment between the two given gambles to receive a return of 14% if a two is thrown?"

The three main conditions

Participants were asked to complete the study in three phases: the *portfolio* condition (*P*), followed by the *allocation* condition (*A*), followed by the *hybrid* condition (*H*). In each of these phases participants were asked to make portfolio selections for various sets of base lotteries using different elicitation procedures. In the portfolio condition P, we asked participants to choose one portfolio from a set of all possible A/B-mixtures (with returns reported to the nearest 0.1%) by turning a dial. We did not explicitly inform them that the options were all combinations of a given pair of base prospects. In the allocation condition A, we asked participants to explicitly allocate their endowment among two risky prospects by clicking and dragging a slider on a 0-100 scale or by entering a number from 0 to 100%. Finally, in the hybrid condition H, we asked participants to allocate their endowment among risky prospects as in condition P, and we also told them the probability distribution of returns that this allocation implied. Further details concerning the elicitation method are provided in the Electronic Companion Paper (ECP).

Convex Treatments

Within each condition (*P*, *A*, and *H*) participants were confronted with multiple trials, each consisting of a different base pair of prospects presented in a random order. Some pairs were modified versions of other pairs. We did this so that we could examine the bias toward even allocation when a similar set of portfolios was made available to participants in different trials but these portfolios entailed a different mixture of base prospects. To understand our notation and how we constructed these modified base pairs, consider Figure 5. Note that any two base prospects, A and B can be depicted as ends of a continuum, and all convex combinations of these prospects (i.e. allocations between A and B) are points along this continuum. Let β denote the portion of the investment that is allocated to prospect B (note that β naturally ranges from 0% on the left to 100% on the right).

For each of the primary base pairs (A,B) from Figure 4, we derived two additional modified base pairs (A’, B) and (A, B’) where A’ = 0.75A + 0.25B and B’ = 0.25A + 0.75B, as depicted in Figure 6. Using the above convention, the prospects A, A’, B’, and B correspond to β = 0%, β = 25%, β = 75% and β = 100%, respectively. We use these values as subscripts to refer to the specific modifications of the primary base pairs or “scenarios” depicted in Figure 4. The notation ${}_0P_{100}$ indicates that in the portfolio elicitation condition the base gambles A and B were used. Similarly, ${}_0P_{75}$ and ${}_{25}P_{100}$ refer to the portfolio elicitation for modified pairs (A,B’) and (A’,B), respectively. The allocation tasks ${}_0A_{100}$, ${}_0A_{75}$, and ${}_{25}A_{100}$ are defined analogously. For instance, for Scenario 3 prospect B’ offers a gain of 17% if the die lands 1 or 2 (=.25*32% + .75*12%), 9% if the die lands 3 or 4, and -10% if the die lands 5 or 6.

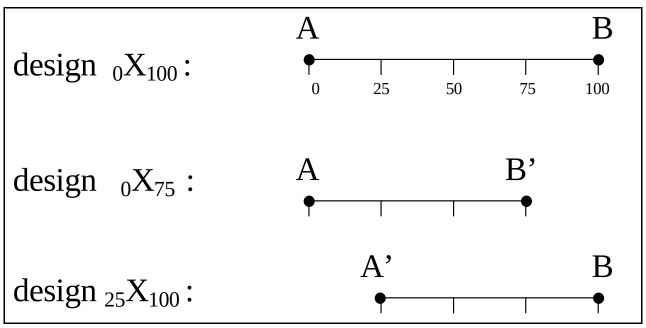

Figure 5: The convex treatments with X = P, A.

Sequence of Trials

For each participant, two of the primary base prospect pairs from Figure 4 were randomly chosen and used to construct 14 specific investment decision trials. In the first phase, there were 6 investment decisions in the portfolio elicitation mode (the designs ${}_0P_{100}$, ${}_0P_{75}$, and ${}_{25}P_{100}$ for both base gamble pairs).[6] In the second phase, there were 6 investment decisions in the allocation mode (${}_0A_{100}$, ${}_0A_{75}$, and ${}_{25}A_{100}$ for both base gamble pairs). Finally, there were two more investment decisions in the third phase, with allocations only for the base pairs (A, B) in the hybrid condition H.[7] Within each phase the ordering of tasks was randomized.

[6] Regrettably, due to a programming bug only five of the data points in treatment P were correctly stored in the database for each participant. As the tasks were randomly ordered within the treatment P, the lost data was evenly distributed over all pairs of prospects and thus about 5/6 of the data in each treatment P category was available.

[7] The hybrid condition was restricted to the base situations ${}_0H_{100}$ to satisfy time constraints.

We told participants that we would pick one participant at random and honor one of that person's choices for real money. We would throw a die in their presence to determine which state of the world occurred, calculate the resulting return of $1,000 invested as the participant designated, and add this sum to a base payment of $250 (the base payment ensured that no participant could actually lose real money). For instance, if an allocation decision from Scenario 3 in Figure 4 was picked and the die landed "4," the total payment could range from $10 (=$250-$1000*24%) if all money was invested into asset A, to $450 (=$250+$1000*20%) if all money was invested into asset B.

<u>Hypotheses and Results</u>

Allocation versus portfolio choice.

Our first hypothesis refers to the manipulation of elicitation mode (portfolio choice, P, versus explicit allocation, A). The present account suggests that a bias toward even allocation will be observed in condition A but be less pronounced (or nonexistent) in condition P. Thus, we assume that implicit allocations in condition P better reflect a participant's "true" unadulterated preference, whereas explicit allocations in condition A should be closer to 50/50, on average.

<u>Hypothesis 3A</u>: The parameter β (the allocation to asset B) should be closer to 50%, on average, for treatment A than treatment P.

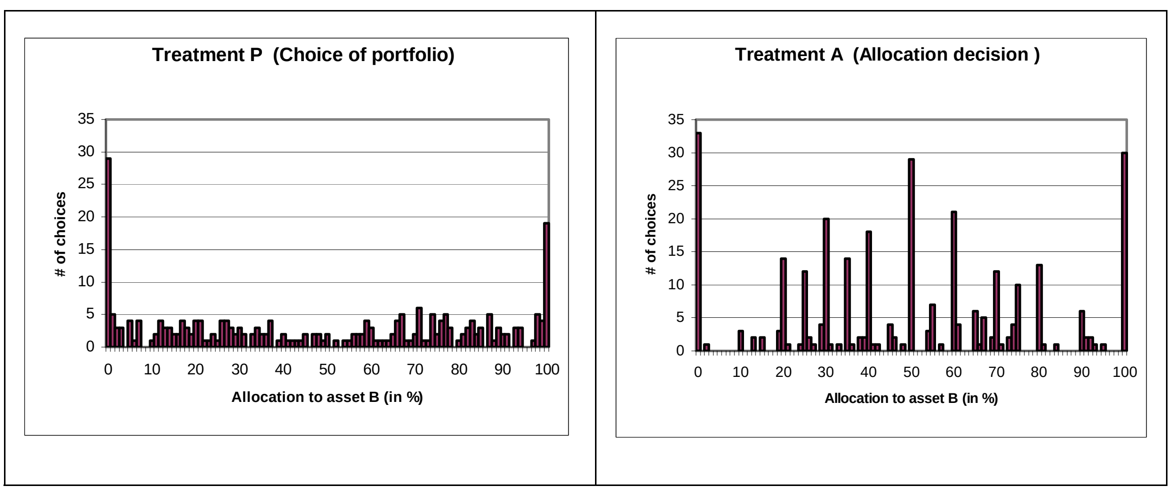

Figure 6: Allocations to asset B for all trials.

Figure 6 displays the distribution of allocations for conditions P and A, for all participants and trials. In condition P, there are two focal points, the extreme allocations 0% and 100% that were chosen in 11.1%

and 7.3% of all cases (n = 260), respectively. There were very few 50/50 allocations (0.77%) in condition P. In contrast, in the allocation condition A (n = 312) the even split (9.3%) is almost as prominent as the extreme allocations 0% and 100% (10.6% and 9.6%, respectively). In general, there are more allocations closer to the middle of the scale.

To explicitly test Hypothesis 3A, we computed the mean absolute distance (MAD) of the stated β from 50% separately for conditions P and A for each participant. Table 4 summarizes the results.[8] Of the 52 participants, 39 (75%) had a lower MAD value in condition A than in treatment P ($p < .0005$ by sign test). The median MAD was larger in the neutral condition P than in the allocation condition A ($p < .001$ by Wilcoxon signed-rank test).[9]

All tasks	Condition P	Condition A
# of subjects with lower MAD value in this treatment (n=52) (One subject had identical MAD values)	12	39
Median MAD	29.4%	21.4%
Stdev. of MAD’s	8.6%	12.1%
	p<0.001	

Table 4: The impact of experimental condition on the MAD, the mean absolute distance of β from 50%.

The role of information concerning portfolio returns.

In order to investigate the extent to which naïve diversification is driven by the elicitation task and to what extent it is driven by the information concerning returns of the aggregate portfolio we added the hybrid condition H in which we asked for allocations but also provided information concerning the probability distribution over outcomes of the aggregate portfolio. We expected that this information would reduce the bias toward even allocation because it would sharpen preferences among aggregate

[8] To base the mean values on exactly the same investment situations, we eliminated the data points from treatment A that correspond to the data missing for treatment P due to the technical problem. Hence, each individual MAD value was generated as the average of five β-distances.

[9] Doing the same analysis separately for each scenario shows that the gamble pairs do not have a strong impact on the effect. The MAD values are generally higher in Scenario 2 (medians 37.5% vs 29.2%) and generally lower in Scenario 3 (23.3% vs. 18.6%), but the differences between treatments are about the same in all scenarios.

portfolios, but it would not eliminate the bias which we expected to be driven in part by the compatibility of the naïve diversification strategy with the task of explicitly allocating money.

<u>Hypothesis 3B</u>: The distance of β from 50% in the hybrid condition H will lie between the corresponding values for conditions P and A, so that providing information concerning returns will reduce the tendency toward naïve diversification but not eliminate it.

The graphical display of allocation decisions in condition H gives some support to Hypothesis 3B (see Figure 7). The allocation pattern seems to be a mixture of the respective patterns for treatment P and treatment A (see Figure 6). For a more quantitative test, we refer again to the mean absolute distances (MAD's) from 50% that are presented in Table 4. Because we only elicited ${}_0H_{100}$ for the hybrid trials, we report MAD values for the ${}_0P_{100}$, ${}_0A_{100}$ and ${}_0H_{100}$ trials only.[10] It can be seen that the ordering of median MAD values for the different conditions is consistent with Hypothesis 3B. The median MAD for the hybrid condition H is 25.0% and thus lies between the median MAD's for condition P (29.75%) and condition A (22.5%). However, due to the limited number of data points and high variance of responses, the differences between condition H and the other two conditions does not achieve statistical significance. Thus, the conclusion that providing distributional information reduces but does not eliminate naïve diversification should be regarded as preliminary.

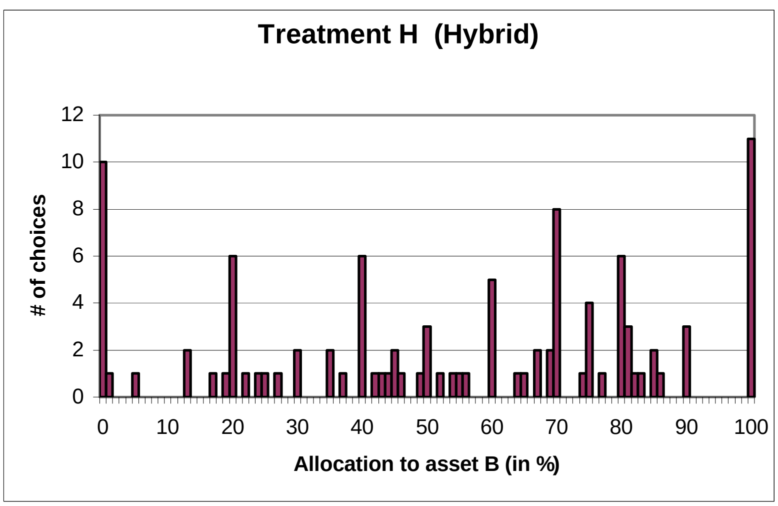

Figure 7: Allocations to asset B for all scenarios in treatment H.

[10] We further eliminated the data points from treatments A and H that were missing in treatment P. In these cases (about 1/6 of the respondents) the MAD actually reflected just one β value.

${}_0X_{100}$ tasks only	Treatment P	Treatment H	Treatment A
Median MAD	29.8%	25.0%	22.5%
Stdev of MAD’s	13.1%	14.5%	14.4%

Table 5: The impact of treatments on the MAD, the mean absolute distance of β from 0.5.

The impact of convex combinations.

Thus far we have provided evidence of procedure dependence: decision maker’s allocation between two risky prospects is influenced by the nature of the elicitation procedure, with a greater bias toward even mixture in allocation than in portfolio choice. We next turn to the manipulation in which we replaced one prospect with a convex combination of the primary base prospects. This allows us to investigate how people behave when they have an opportunity to replicate their original preference by allocating between a modified set of base prospects and also provides a second test of procedure dependence.

Because allocations among both the base prospects and modified base prospects presumably reflect a 50-50 bias it is somewhat complicated to make a prediction concerning the direction of the manipulation effect. Specific properties of the modification determine whether the β stated by a biased person for the modified pair will be higher or lower than the $\beta^{unbiased}$ that would replicate the choice of the original task. A detailed description of these properties and a formal derivation of the predictions are presented in the electronic companion paper (ECP). For the two specific modifications that we used in our experimental analysis, ${}_0A_{75}$ and ${}_{25}A_{100}$, the following hypotheses can be stated as special cases of the general result:

Hypothesis 3C: When participants allocating among modified base prospects (i.e. in situation ${}_0A_{75}$ or ${}_{25}A_{100}$) can replicate their allocation from ${}_0A_{100}$ by choosing some $\beta^{unbiased}$, the actual allocations β in the modified pair will be

- lower than $\beta^{unbiased}$ for the modification ${}_0A_{75}$ and
- higher than $\beta^{unbiased}$ for the modification ${}_{25}A_{100}$.

This effect will be less pronounced in direct portfolio choice.

Note that if an allocation β > 75% was chosen in task ${}_0A_{100}$, the same portfolio was not available in task ${}_0A_{75}$. Likewise, if an allocation β < 25% was chosen in task ${}_0A_{100}$, the same portfolio was not available in task ${}_{25}A_{100}$. Thus, for the test of Hypothesis 3C we exclude these data points.[11] The results concerning Hypothesis 3C are summarized in the upper part of Table 6.

From the 84 allocation comparisons between tasks ${}_0A_{100}$ and ${}_0A_{75}$ the hypothesized pattern of bias ($\beta < \beta^{unbiased}$) is observed in 51 cases (60.7%), the opposite pattern in only 27 cases (32.1%). The median difference between the observed β and $\beta^{unbiased}$ is –10.0%. Similar results are found for the comparison of task ${}_0A_{100}$ with task ${}_{25}A_{100}$ (56% of cases in line with and 38% opposite to the hypothesis). The median difference between β and $\beta^{unbiased}$ is +7.5%, thus also for ${}_{25}A_{100}$ the effect is in the hypothesized direction. Overall, 98 (58.3%) observations are consistent with Hypothesis 3C and only 59 (35.1%) oppose Hypothesis 3C. Note again, that there would be no systematic difference between observed and unbiased β's if people allocated at random.

	${}_0A_{75}$ vs. ${}_0A_{100}$			${}_{25}A_{100}$ vs. ${}_0A_{100}$			All comparisons		
	$\beta<\beta^{unbiased}$ (hypoth.)	$\beta>\beta^{unbiased}$	No bias	$\beta>\beta^{unbiased}$ (hypoth.)	$\beta<\beta^{unbiased}$	No bias	β in hyp. direction	β not in hyp. direct.	No bias
# of observations	**51** (60.7%)	**27** (32.1%)	**6** (7.2%)	**47** (56.0%)	**32** (38.1%)	**5** (5.9%)	**98** (58.3%)	**59** (35.1%)	**11** (6.6%)
median of $(\beta-\beta_{unbiased})$	-10.0%			+7.5%					
	${}_0P_{75}$ vs. ${}_0P_{100}$			${}_{25}P_{100}$ vs. ${}_0P_{100}$			All comparisons		
	$\beta<\beta^{unbiased}$	$\beta>\beta^{unbiased}$	No bias	$\beta>\beta^{unbiased}$	$\beta<\beta^{unbiased}$	No bias	β in hyp. direction	β not in hyp. direct.	No bias
# of observations	**18** (38.3%)	**28** (59.6%)	**1** (2.1%)	**26** (54.2%)	**21** (43.7%)	**1** (2.1%)	**44** (46.3%)	**49** (51.6%)	**2** (2.1%)
Median of $(\beta-\beta_{unbiased})$	+1.0%			+2.2%					

Table 6: Results of the convex manipulation for treatments A and P.

[11] By excluding these data points we get an unbiased procedure in the sense that no systematic effect would be observed if participants made random choices.

Determining the significance of the observed differences is somewhat tricky. Each subject contibutes up to four data points (β-$\beta^{unbiased}$) to the figures in Table 6.[12] So it seems straightforward to aggregate these data points to get one independent score for each participant. However, it turns out that the obvious aggregation procedure (averaging the differences with reversed sign for ${}_{25}A_{100}$) is systematically biased due to the fact that not all four data points exist for all participants.[13] To avoid this problem we have to look at the treatments separately and to combine just two (or even fewer) data points to receive a participant's average (β-$\beta^{unbiased}$) difference per treatment. Unfortunately this leaves more noise in the data and decreases the statistical significance of our results.[14] For the ${}_{0}A_{75}$ treatment, the median of the participant's average difference (β-$\beta^{unbiased}$) is –12% (number not reported in the table). This difference approaches statistical significance (p=0.06, one-tailed Wilcoxon signed-rank test). For the ${}_{25}A_{100}$ treatment, the median of the participant's average difference (β-$\beta^{unbiased}$) is +2.9% (number not reported in the table). This number is not significantly different from 0.

Consistent with Hypothesis 3C we observed no significant effect of the convex manipulation in the direct portfolio choice (see lower part of Table 6). Fewer than half of the portfolio choices were biased in the direction that was predicted for treatment A. Roughly the same number of choices pointed in the opposite direction. Even though the effect for treatment ${}_{25}P_{100}$ is in the same direction as predicted for

---

[12] These are the differences (β-$\beta^{unbiased}$) for two scenarios and both treatments ${}_{0}A_{75}$ and ${}_{25}A_{100}$. Some of these data points are missing though if the ${}_{0}A_{100}$ allocations were not reproducable in the modified conditions (i.e. if the ${}_{0}A_{100}$ allocation was above 75% or below 25%).

[13] The problem occurs, because in the averaging process (β-$\beta^{unbiased}$) differences get relatively more weight if their counterparts (β-$\beta^{unbiased}$ differences in the other modified condition) are missing. These overweighted data points with missing counterparts, however, are more likely to make just a small contribution in the predicted direction or even point in the opposite direction, because they correspond to specific (extreme) ${}_{0}A_{100}$ allocations. Therefore the procedure would generate a systematic effect even for random choices. To illustrate, consider a situation where the ${}_{0}A_{100}$ allocation is close to 0%. In this case, we have $\beta^{unbiased} \approx 0$ in the ${}_{0}A_{75}$ condition and it is hardly possible to observe an effect of $\beta<\beta^{unbiased}$ as predicted in Hypothesis 3C. If there is some noise in the data, it is likely that we would observe an effect opposite to our prediction. As long as we consider isolated data points this is not a problem, as this distortion is perfectly balanced by opposite distortions for other allocations. However, through the averaging procedure we produce a systematic bias as the above mentioned distortions receive extra weight (note that in our example the counterpart for the ${}_{25}A_{100}$ is missing).

[14] Our data (also the data in the P condition) is in fact surprisingly noisy and could lead to the impression that individuals have no consistent preferences with respect to such lotteries. However, to some extent this might also be attributed to our experimental design in which seemingly different allocations can result in very similar portfolios (e.g. in Scenario 1 presented in Figure 4, the allocation of 60% to prospect B and the allocation of 100% to prospect B result in very similar probability distributions over returns).

treatment A, it turns out to be far from significant when we aggregate the data within subject as we did for treatment A. If we compare the proportion of participants that are biased in the predicted direction in the allocation and portfolio choice, we find a significant difference for treatment ${}_{0}X_{75}$ (64% vs. 37%, $p<0.02$), while the difference (55% vs. 53%) is not significant for treatment ${}_{25}X_{100}$ (aggregated numbers are not reported in the table).

General Discussion

In this paper, we have explored biases that people exhibit when allocating money among various risky or uncertain prospects. In the present studies we found consistent evidence that relatively sophisticated participants apply a naïve diversification heuristic, even when prospects are simple, well-specified lotteries and there are incentive-compatible payoffs. In addition, we provide evidence of three important new implications of naïve diversification: partition dependence, unit dependence, and procedure dependence. We conclude with a more detailed review of our experimental results, a discussion of promising avenues for future research, manifestations of diversification strategies in other domains of judgment and choice, and prescriptive implications of the present work.

In Study 1, we demonstrated that the way in which the set of available prospects is partitioned influences the way in which people allocate money among these prospects. Different partitions of the asset space can be induced by grouping of the available assets hierarchically and using a multistage allocation procedure. We experimentally demonstrated that assets receive a greater investment in the aggregate portfolio if they are presented by themselves rather than grouped with alternative assets. This result was observed for hypothetical investments of familiar asset classes (stocks, bonds, real-estate) that were grouped by vendor, and also for simple, well-specified lotteries that were grouped arbitrarily. This result is of practical relevance as real allocation decisions are commonly guided by a hierarchical structure such as clients to accounts, divisions to a firm,[15] or companies to industries. Investment advisors

[15] For evidence of partition dependence in allocation of capital to divisions within a firm, see Bardolet, Fox & Lovallo (2005).

often ask for intended allocations to general categories such as stocks and bonds before they address the detailed allocations within these categories. Many 403b retirement saving plans (offered by many American universities) require participants to first indicate allocations to vendors, then later to specific funds, separately for each vendor. Obviously, this procedure can strongly influence the allocation of funds to asset classes, especially if different vendors specialize in particular types of products.

In Study 2, we demonstrated that the particular unit in which the allocations are specified influences the nature of the portfolio that is selected. Thus, an even split in terms of the money invested may not correspond to an even split in terms of the number of shares purchased. Our experiments found that participants allocated a larger portion of their endowment to the more expensive asset when indicating the proportion of shares to be purchased for each asset than when indicating the proportion of money to be spent on each asset. Again, this result was observed for well-specified gambles as well as for real stocks. The finding is interesting as the unit in which decision makers are asked to indicate their investment varies in different situations. For instance, stock investments are typically expressed in terms of a number of shares purchased whereas mutual fund investments are typically expressed in monetary terms. Analogously, firms can designate the assignment of people to tasks in terms of number of individuals, hours of labor, or dollars spent, which could have different implications if there are differences in salary or marginal product of labor for different classes of employee or different tasks.

In Study 3, we obtained further evidence of naïve diversification in a tightly controlled environment with well-specified gambles, incentive-compatible payoffs, and a within-subject experimental design. Moreover, we obtained strong evidence of procedure dependence. When participants were asked to allocate their endowment among two simple gambles, the observed allocations were on average closer to 50/50 than when they were asked to choose from a set of possible portfolios that could be obtained from a blend of base prospects. Providing information concerning the probability distribution over returns of possible allocations seemed to moderate, but not eliminate, the tendency toward an even split. Moreover, when we modified one of the base lotteries, explicit allocations exhibited a stronger bias toward 50-50 than were choices among possible portfolios. Thus, the tendency to split money evenly among the available options seems to be driven largely by the task of allocating rather than by other factors such as the availability of information on aggregate returns or expected value, and it extends beyond the domain of vaguely specified personal investments.[16]

Other manifestations of diversification and partition dependence

We assert that naïve diversification may be an instance of a more pervasive cognitive strategy that people employ across a wide range of decision making and judgment tasks: When allocating a scarce resource over a fixed set of options, people typically rely on an even allocation across all groups into which the option set is partitioned, adjusting this distribution to the extent that they can distinguish among options. These even allocation strategies might be thought of as "maximum entropy heuristics." In the case of *naïve diversification*, the scarce resource is money and the set of options is risky or uncertain prospects. Similarly, research in consumer choice has revealed a tendency to seek variety over consumer products to be consumed in the future (Simonson, 1990; Read & Loewenstein, 1995; Fox, Ratner & Lieb, in press). In this case, the scarce resource is a fixed number of selections to be made and the set of options may be candies that are naturally partitioned by variety. Research in distributive justice has shown that people often rely on an *equality heuristic:* they tend to allocate resources or responsibilities evenly among group members (Harris & Joyce, 1980; Messick, 1993; Roch, et al., 2000). In this case the scarce resources are benefits and burdens and the set of options is individuals or groups. In judgment under uncertainty, people seem to rely on an intuitive application of the *principle of insufficient reason*, allocating probabilistic belief evenly across all events that might occur (Fox & Rottenstreich, 2003; Fox & Levav, 2004; Fox & Clemen, in press). In this case the scarce resource is probabilistic belief and the set of options is events into which the sample space is partitioned. In multiattribute choice, people exhibit *attribute splitting effects* in which they are biased toward assigning equal weight to each of the features

[16] The within-subject experimental design of Study 3 provided us with a further insight into how individuals allocate money among risky prospects. In general, individuals do not seem to have particularly consistent preferences over return distributions, even in the simple scenarios we presented. When we provided participants with an opportunity to indicate their most preferred portfolio using different elicitation modes or slightly modified base lotteries, they exhibited only a small amount of consistency in their choices.

that are used to characterize options so that features receive greater weight when they are split into subfeatures (Weber, Eisenführ & von Winterfeldt, 1988). In this case the scarce resource is attribute weight and the set of options is attributes.

Different maximum entropy heuristics may be motivated by different concerns (e.g., desire to minimize risk, to be open to new experiences, to maintain harmony of the group, to minimize error) and the extent to which people rely on them may vary with different factors (e.g., levels of knowledge or sophistication, cognitive effort, etc.). In all cases, use of the heuristic will give rise to some of the systematic dependencies that we have observed in this paper (for a review see Fox, Bardolet & Lieb, 2005).

Future Research

There are many promising avenues for future research that build on the present investigation. Notably, one might explore the factors that moderate naïve diversification and partition dependence, such as increasing knowledge. One might have expected a priori that knowledge concerning financial markets, the basic principles of diversification, and statistics could inoculate participants against naïve diversification. On the other hand, we recruited sophisticated participants in most of our studies, most of them MBA or graduate students, some even specializing in banking and finance, and nevertheless found strong biases.[17] This suggests that providing investors with statistical and financial background knowledge would probably not be sufficient to inoculate them against naive diversification and the associated dependencies. Future research might consider the question of what specific type of knowledge or information would help investors avoid these pitfalls. Does knowledge concerning the biases help? To what extent does substantive knowledge concerning the potential investments play a role? In addition, one might examine the role of other potential moderators. For instance, do investors who are under time pressure or cognitive load exhibit more pronounced biases? To what extent will naïve diversification be exacerbated or mitigated when decisions are made by groups rather than individuals?

Prescriptions

---

[17] In Study 1A, we asked for self-rated knowledge and found no significant effect of knowledge on the strength of the bias.

Future research on naïve diversification and associated biases may help practitioners design asset allocation procedures that are less susceptible to bias than those currently in use. Based on what we have discovered thus far we suggest some preliminary prescriptions. First, practitioners should take special care not to bias respondents through the hierarchical grouping of particular investments. Potential investments should either be presented non-hierarchically, or should be grouped into classes of instruments whose returns are highly correlated (e.g., equities, annuities, commodities). Second, care should be taken in choosing the units to be allocated so that the units allocated (e.g., dollars) are compatible with the units that are valued by the decision maker (e.g., future dollar value of the portfolio). Third, practitioners might search for an elicitation mode that minimizes bias, provided that this mode represents a reliable and valid measure of preferences. Fourth, practitioners might present participants with elicitations that rely on different partitions, units, and procedures, and provide participants an opportunity to resolve for themselves any inconsistencies that might arise.

A second approach to debiasing is to make the task easier for decision makers. This might be accomplished by educating them concerning the potential investments and the correlation of their returns, providing decision support tools such as decision trees and Monte Carlo simulations. Alternatively, practitioners might provide a simplified menu of representative investments.

Finally, a more paternalistic approach to debiasing would be to design an allocation procedure in a way that is expected to satisfy the investor's best interests and avoid any biases. For instance, portfolios could be constructed and periodically rebalanced automatically in ways that are designed to maximize the expected utility of the respondent, based on an independent assessment of the respondent's utility function and a statistical analysis of the expected returns and correlations among the available investments (For a promising early attempt see Goldstein, Johnson & Sharpe, 2005).

References

Bardolet, D., Fox, C.R., & Lovallo, D. (2005). Partition dependence in organizational capital allocation: An experimental investigation. Working paper, UCLA Anderson School.

Benartzi, S. & Thaler, R. (2001): Naïve diversification strategies in defined contribution saving plans. *American Economic Review*, 91, 79-98.

Benartzi, S. & Thaler, R. (2002): How Much Is Investor Autonomy Worth? *Journal of Finance*, 57, 1593-1616.

Fischer, G.W. & Hawkins, S.A. (1993). Scale Compatibility, Strategy Compatibility and the Prominence Effect. *Journal of Experimental Psychology: Human Perception and Performance*, 19, 580-597.

Employee Benefit Research Institute (2004). Issue Brief #272. http://www.ebri.org/ibpdfs/0804ib.pdf.

Fox, C.R., Bardolet, D., & Lieb, D. (2005). Partition dependence in decision analysis, managerial decision making, and consumer choice. Chapter to appear in Zwick, R. & Rapoport, A. *Experimental Business Research, Volume III* (Dordrecht, The Netherlands: Kluwer).

Fox, C.R. & Clemen, R.T. (in press): Assigning subjective probabilities to event trees: Partition dependence and bias toward the ignorance prior. *Management Science*, forthcoming.

Fox, C.R. & Levav, J. (2004). Partition-edit-count: Naïve extensional reasoning in conditional probability judgment. *Journal of Experimental Psychology: General, 133*, 626-642.

Fox, C.R. & Rottenstreich, Y. (2003): Partition priming in judgment under uncertainty. *Psychological Science,* 14, 195-200.

Fox, C.R., Ratner, R.K. & Lieb, D. (in press): How Subjective Grouping of Options Influences Choice and Allocation: Diversification Bias and the Phenomenon of Partition Dependence. *Journal of Experimental Psychology: General,* forthcoming.

Goldstein, D. G., Johnson, E. J., & Sharpe, W. F. (2005). Measuring consumer risk-return tradeoffs. Manuscript submitted for publication.

Harris, R. J., & Joyce, M. A. (1980). What's fair?: It depends on how you phrase the question. *Journal of Personality and Social Psychology*, 38, 165-179.

Messick, D. (1993): Equality as a decision heuristic. In B.A. Mellers & J. Baron (Eds.) Psychological perspective on justice (pp. 11-31). Cambridge: Cambridge University Press.

Poterba, J.M., Venti, S.F., & Wise, D.A. (2001): The Transition to Personal Accounts and Increasing Retirement Wealth: Macro and Micro Evidence. NBER Working Paper 8610.

Read, D. & Loewenstein, G. (1995): The diversification bias: Explaining the difference between prospective and real-time taste for variety. *Journal of Experimental Psychology: Applied,* 1, 34-49.

Roch, S., Samuelson, C., Allison, S., & Dent, J. (2000). Cognitive load and the equality heuristic: A two stage model of resource overconsumption in small groups. *Organizational Behavior and Human Decision Processes,* 83, 185-212.

Simonson, I. (1990): The Effect of Purchase Quantity and Timing on Variety-Seeking Behavior. *Journal of Marketing Research*, 27, 150-162

Slovic, P., Griffin, D. & Tversky, A. (1990): Compatibility Effects in Judgment and Choice. in: Hogarth, R.M. (Ed.): Insights in Decision Making: A Tribute to Hillel J. Einhorn. University of Chicago Press, 5-27.

Weber, M., Eisenführ, F. & von Winterfeldt, D. (1988): The Effect of Splitting Attributes in Multiattribute Utility Measurement. *Management Science*, 34, 431-445.

Electronic Companion Paper
Bias in allocation under risk and uncertainty:
Partition dependence, unit dependence, and procedure dependence

I. Methodological Details of Study 3

The upper-left section of Figure 1 displays a screenshot from the portfolio condition P. We asked participants to use a dial to explore different portfolios in the choice set. Specific dial positions corresponded to specific mixtures of the (hidden) base prospects with dial positions continuously mapped onto the different allocations.[18] When participants began a new trial in condition *P*, no portfolio outcomes were displayed nor was the position marker visible until they first clicked on the dial.[19] Participants then could click and drag the dial with the computer mouse or use specific keys to make large or fine adjustments. To facilitate this process a tentative choice could be temporarily stored for comparison with other options. In order to simplify the task, the probability distribution of returns was always displayed with the highest possible return appearing at the top and the lowest possible return at the bottom.

Once participants indicated their final selection, they were prompted to consider this selection against similar portfolios that were also available in the offered set. Participants were then offered an opportunity to reconsider their final choice if they wished to do so. This procedure is illustrated in the upper-right section of Figure 1.

In condition A participants were asked to explicitly allocate their endowment among the base gambles. Both base gambles were displayed on the screen and allocations could be specified by either typing numbers into a text field (the percentage of the endowment to be allocated to the associated base gamble) or clicking and dragging a visual slider (see the lower-left section of Figure 1). The text field and slider were synchronized and text entries were forced to be in the [0,100] interval. To avoid any

[18] The two extreme 100% allocations were located on opposite sides of the dial. Turning the dial either side would continuously change the allocation from the one extreme to the other. Thus, by a 360 degree turn of the dial each allocation (other than the extremes) would appear exactly twice.

unintended anchoring, no starting allocation was specified when participants began a new task; that is, text fields were empty and there was no marker on the slider bar until participants first clicked onto the bar or typed a number in the text field. We randomized the side of the screen on which each gamble appeared for each trial and each participant.

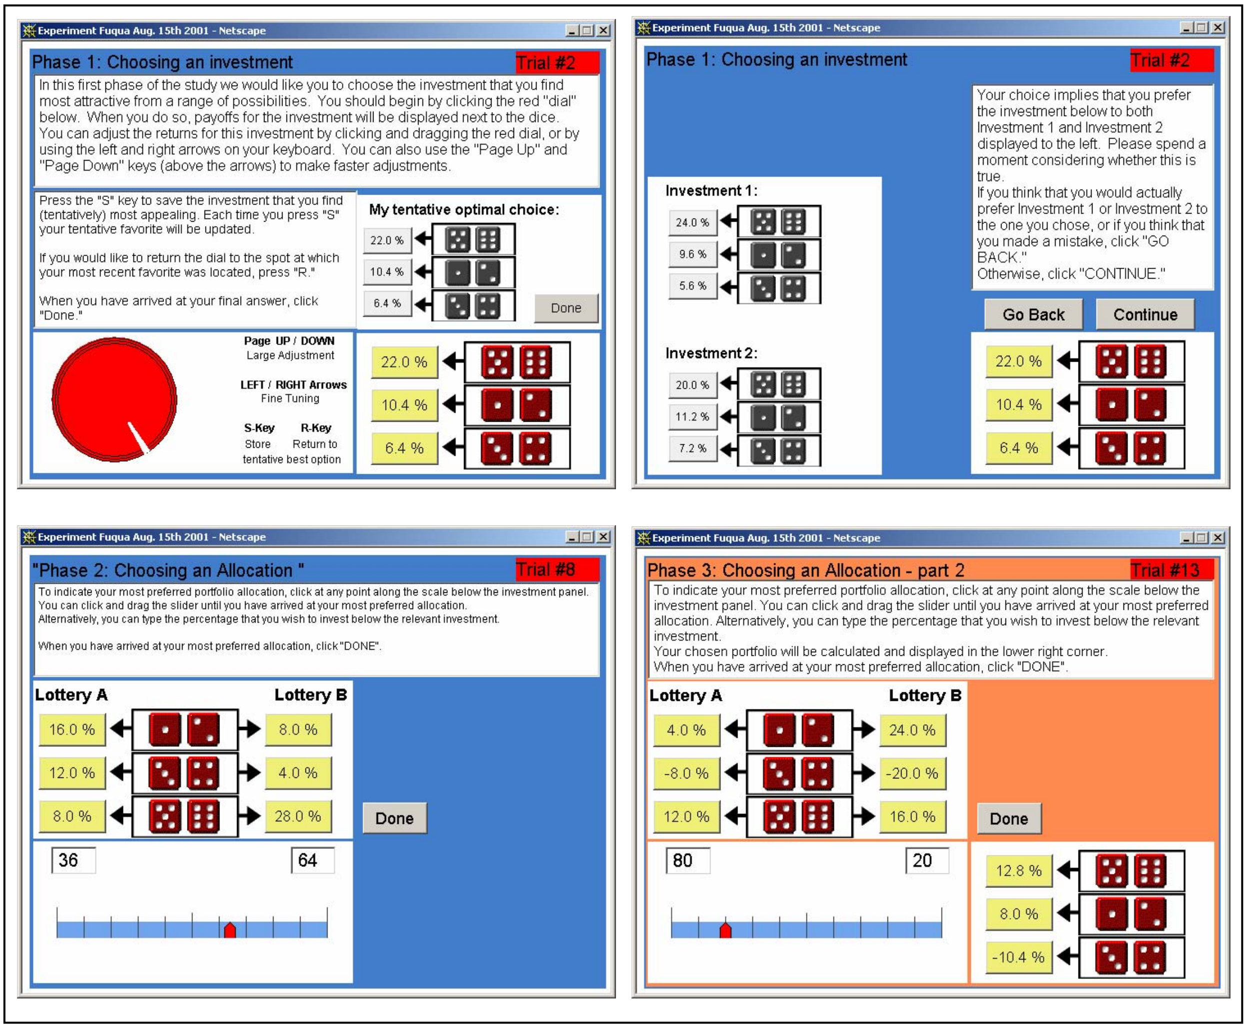

Figure 1: Screenshots from Study 3.

In the hybrid condition H, the same elicitation mode as in condition A was used, but in addition, information concerning the probability distribution of aggregate portfolio returns was presented as in condition P. The slider was synchronized with the portfolio return display, such that for each temporarily

[19] This was done to avoid an anchoring on some specific starting allocation. Participants themselves generated the dial marker at the position on the dial that they first clicked. The exact mapping from dial position to portfolio allocation was also randomized for each participant and task.

chosen allocation, participants got immediate information about the implied return distribution (see the lower-right section of Figure 1).

II. Formal Characterization of Hypothesis 3c

To explain our analysis of convex treatments we refer to Figure 2. The top line, ${}_0A_{100}$, depicts allocations among base prospects A and B from 0% to 100% allocated in prospect B. The bottom line, ${}_0A_{75}$, depicts allocations among base prospect A and modified base prospect B' which is itself a mixture of 25% A and 75% B.[20] This line is $\frac{3}{4}$ the length of the top line. Each mixture of the original base prospects A and B with $\beta \leq 75\%$ can be reproduced by a mixture of the modified prospects A and B' on the compressed scale.

Now, take some stated allocation ${}_0\beta_{100}^{\text{stated}}$ between the primary base prospects A and B and consider a consistent decision maker that exhibits no naïve diversification. She should report ${}_0\beta_{75}^{\text{unbiased}}$ for the modified task which is the point directly below on the compressed scale. We would thus observe ${}_0\beta_{75}^{\text{unbiased}} = \frac{4}{3} \cdot {}_0\beta_{100}^{\text{stated}}$. Now let us suppose instead that an individual exhibits a bias of naïve diversification, such that the "true" underlying preference, ${}_0\beta_{100}^{\text{true}}$, is distorted toward 0.5 to result in ${}_0\beta_{100}^{\text{stated}}$. We can formalize the concept by introducing a parameter λ ($0 \leq \lambda \leq 1$) that represents the degree of bias toward the middle of the scale, i.e. ${}_0\beta_{100}^{\text{stated}} = \lambda \cdot \frac{1}{2} + (1-\lambda) \cdot {}_0\beta_{100}^{\text{true}}$. Thus, for $\lambda = 0$ an individual is completely unbiased, and for $\lambda = 1$ she relies entirely on naïve diversification.

[20] To distinguish the different βs, we use the notation ${}_0\beta_{100}$ for mixtures of A and B and the notation ${}_0\beta_{75}$ for mixtures of A and B'.

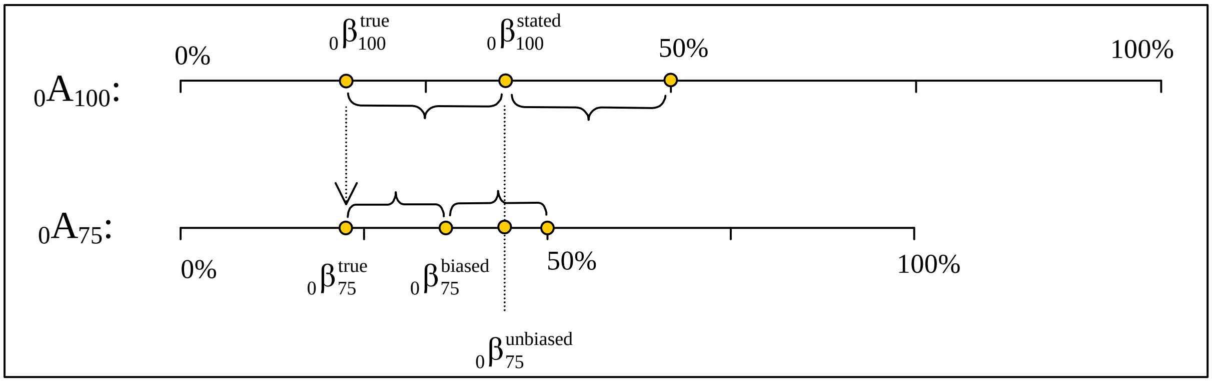

Figure 2: An illustration of predictions concerning convex combinations.

Now suppose in our hypothetical example there is a moderate degree of naïve diversification, e.g. $\lambda = 0.5$. In this case we can derive the "true" portfolio preference $_0\beta_{100}^{true}$ from the observed choice $_0\beta_{100}^{stated}$, map this onto a "true" portfolio preference for the modified pair of based prospects, $_0\beta_{75}^{true}$, and predict that the stated probability $_0\beta_{75}^{stated}$ for this individual is $_0\beta_{75}^{biased}$. The argument is graphically presented in Figure 2. It can be seen that in this specific situation $_0\beta_{75}^{biased}$ is smaller than $_0\beta_{75}^{unbiased}$, a relation that will obviously hold for all other λ smaller than 1, too. In fact, with a straightforward extension of our notation to the set of all convex manipulations $_fX_g$ (with $0 \le f < g \le 100$) a more general proposition can be shown:[21]

$$\text{For each } \lambda \in [0,1): \quad {}_f\beta_g^{biased} < (>=)\ {}_f\beta_g^{unbiased} \quad \Leftrightarrow \quad \frac{f+g}{2} < (>=)\ 50$$

As two special cases of this general proposition we get:

$$_0\beta_{75}^{biased} < {}_0\beta_{75}^{unbiased} \quad \text{and} \quad {}_{25}\beta_{100}^{biased} > {}_{25}\beta_{100}^{unbiased}\ .$$

That is, if people rely at all on naïve diversification we should observe the stated allocation $_0\beta_{75}^{biased}$ for the modified pair (A, B') to be less than four-thirds the allocation $_0\beta_{100}^{stated}$ for the

[21] A proof of this proposition is available from the authors upon request.

original base pair (because ${}_{0}\beta_{75}^{\text{unbiased}} = \frac{4}{3} \cdot {}_{0}\beta_{100}^{\text{stated}}$). And we should observe a stated allocation ${}_{25}\beta_{100}^{\text{biased}}$ for the modified pair (A', B) that is higher than $\frac{4}{3} \cdot {}_{0}\beta_{100}^{\text{stated}} - \frac{1}{3}$ ($= {}_{25}\beta_{100}^{\text{unbiased}}$).

This leads to a more formal characterization of our research hypothesis:

Hypothesis 3C: ${}_{0}\beta_{75}^{\text{stated}} < {}_{0}\beta_{75}^{\text{unbiased}}$ ($= \frac{4}{3} \cdot {}_{0}\beta_{100}^{\text{stated}}$) and ${}_{25}\beta_{100}^{\text{stated}} > {}_{25}\beta_{100}^{\text{unbiased}}$ ($= \frac{4}{3} \cdot {}_{0}\beta_{100}^{\text{stated}} - \frac{1}{3}$) for the explicit allocation decisions in condition A, and this tendency will be less pronounced for the portfolio choices in condition P.